arXiv:1704.06361v1 [cs.DM] 20 Apr 2017

Shared processor scheduling

Dariusz Dereniowski
Faculty of Electronics, Telecommunications and Informatics, Gdańsk University of Technology, Gdańsk, Poland

Wiesław Kubiak
Faculty of Business Administration, Memorial University, St. John's, Canada

Abstract

We study the shared processor scheduling problem with a single shared processor where a unit time saving (weight) obtained by processing a job on the shared processor depends on the job. A polynomial-time optimization algorithm has been given for the problem with equal weights in the literature. This paper extends that result by showing an $O(n \log n)$ optimization algorithm for a class of instances in which non-decreasing order of jobs with respect to processing times provides a non-increasing order with respect to weights — this instance generalizes the unweighted case of the problem. This algorithm also leads to a $\frac{1}{2}$-approximation algorithm for the general weighted problem. The complexity of the weighted problem remains open.

Keywords: divisible jobs, scheduling, shared processor

1 Introduction

Consider a subcontracting system in which each agent j has a job of duration p_j to be executed. Such an agent can perform the work by itself, in which case the job ends after p_j units of time, or it can send (subcontract) a part of length $s_j \leq p_j/2$ of this job to a subcontractor for processing. The subcontractor needs to complete this piece of agent's j job by $p_j - s_j$, i.e., the speedup in terms of the completion time that the agent achieves in this scenario is exactly s_j, or in other words, the work of agent j is completed at time moment $p_j - s_j$. Whenever $s_j > 0$, the subcontractor is rewarded by agent j: the payoff of executing s_j units of j-th agent's job is $s_j w_j$. The goal of the subcontractor is to maximize its total payoff under the condition that the parts of jobs received from different agents cannot be executed simultaneously by the subcontractor. Thus, in this subcontracting system all agents try to minimize completion times of their jobs (the parameters p_j and w_j are fixed for each agent j), i.e. they are willing to commission the biggest possible part of their jobs to the subcontractor. The subcontractor is the party that decides what amount, if any, of each job to process in order to maximize its total payoff.

The shared processor scheduling problem can be placed in a wider context of scheduling with presence of private (local) processors (machines), available only to a particular job or a set of jobs, and shared (global) processors that are available to all jobs. Then, some additional rules are given in order to specify the conditions under which a job can gain access to a shared processor in such systems. These systems can be run as either centralized or decentralized. The former typically has a single optimization criterion forcing all

Emails: deren@eti.pg.gda.pl (Dariusz Dereniowski) and wkubiak@mun.ca (Wiesław Kubiak)

parties to achieve the same goal. The latter emphasizes that each party is trying to optimize its own goal, which may (and often does) lead to problems having no solutions which are optimal for each agent (job) individually. These problems can be seen as multi-criteria optimization or coordination problems. The latter and can be further subdivided into problems in which agents have complete knowledge about resources of other agents (complete information games) and problems without such a full knowledge (distributed systems) in the search for coordinating mechanisms. This work falls into the category of centralized problems as the subcontractor is deciding on the schedule that reflects its best interest.

The outline of this paper is as follows. In the next section we briefly survey the related work to provide a state of the art overview. Section 3 gives a formal statement of the scheduling problem we study and it introduces the necessary notation. Then, in Section 4, we recall some facts related to the problem, mainly the fact that when computing optimal schedules one may restrict attention to schedules that are called *synchronized*. This generally greatly reduces the formal arguments and algorithmic approach. Section 5 considers a restricted version of the problem in which it is assumed that for any pair of jobs, neither of the jobs can have weight and processing time to be strictly smaller than the other. We give an $O(n \log n)$-time optimization algorithm for this case, and we use it subsequently as a building block to obtain an $O(n \log n)$-time $1/2$-approximation algorithm for the general case in Section 6.

2 Related work

The shared processor scheduling problem has recently been studied by Vairaktarakis and Aydinliyim [10], Hezarkhani and Kubiak [7], and Dereniowski and Kubiak [6]. Vairaktarakis and Aydinliyim [10] consider the (unweighted) problem with a single shared processor and with each job allowed to use at most one time interval on the shared processor. This case is sometimes referred to as non-preemptive since jobs are not allowed preemption on the shared processor. [10] proves that there are optimal schedules that complete job execution on private and shared processor at the same time, we call such schedules *synchronized*, for the non-preemptive case with equal weights. It further shows that this guarantees that sequencing jobs in ascending order of their processing times leads to an optimal solution for the case. [7] observes that this algorithm also gives optimal solutions to the preemptive problem, where more than one interval can be used by a job on the shared processor, provided that all weights are equal. [6] consider shared multi-processor problem proving its strong NP-hardness and giving an efficient, polynomial-time algorithm for the shared multi-processor problem with equal weights. Also, it is shown in [6] that synchronized optimal schedules always exist for weighted multi-processor instances. Vairaktarakis and Aydinliyim [10] , Vairaktarakis [9], and Hezarkhani and Kubiak [7] also study decentralized subcontracting systems focusing on coordinating mechanisms to ensure their efficiency.

The motivation to study the shared processor scheduling problem comes from diverse applications. Vairaktarakis and Aydinliyim [10] consider it in the context of supply chains were subcontracting allows jobs to reduce their completion times by using a shared subcontractor's processor. Bharadwaj et. al. [3] use the divisible load scheduling to reduce a job completion time in parallel and distributed computer systems, and Anderson [2] argues for using batches of potentially infinitely small items that can be processed independently of other items of the batch in scheduling job-shops. We refer the reader to Dereniowski and Kubiak [6] for more details on these applications.

We also remark multi-agent scheduling models in which each agent has its own optimality criterion and performs actions aimed at optimizing it. In these models, being examples of decentralized systems, agents usually have a number of non-divisible jobs to execute (depending on the optimization criterion this may be seen as having one divisible job, but restricted by allowing preemptions only at certain specified points).

For minimization of weighted total completion time in such models see Lee et. al. [8] and weighted number of tardy jobs see Cheng, Ng and Yuan [5]. Bukchin and Hanany [4] give an example of a game-theoretic analysis to a problem of this type. For overviews and further references on the multi-agent scheduling we refer to the book by Agnetis et. al. [1].

3 Problem formulation

We are given a set $\mathcal{J}$ of n preemptive jobs. Each job $j \in \mathcal{J}$ has its processing time p_j and weight w_j. With each job $j \in \mathcal{J}$ we associate its *private* processor denoted by $\mathcal{P}_j$. Moreover, there exists a single shared processor, denoted by $\mathcal{M}$, that is available for all jobs. We follow the convention and notation from [6] to formulate the problem in this paper.

A schedule $\mathcal{S}$ is *feasible* if satisfies the following conditions:

- each job $j \in \mathcal{J}$ executes non-preemptively in a *single* time interval $(0, C_{\mathcal{S}}^{\mathcal{P}}(j))$ on its private processor and there is a (possibly empty) collection of open intervals $\mathcal{I}_j$ such that j executes non-preemptively in each time interval $I \in \mathcal{I}_j$ on the shared processor,
- for each job $j \in \mathcal{J}$,
$$C_{\mathcal{S}}^{\mathcal{P}}(j) + \bigcup_{I \in \mathcal{I}_j} |I| = p_j,$$
- the time intervals in $\bigcup_{j \in \mathcal{J}} \mathcal{I}_j$ are pairwise disjoint.

Given a feasible schedule $\mathcal{S}$, for each job $j \in \mathcal{J}$ we call any maximal time interval in which j executes on both private $\mathcal{P}_j$ and shared $\mathcal{M}$ simultaneously an *overlap*. The total overlap t_j of job j equals the sum of lengths of all overlaps for j. The *total weighted overlap* of $\mathcal{S}$ equals

$$\Sigma(\mathcal{S}) = \sum_{j \in \mathcal{J}} t_j w_j.$$

A feasible schedule that maximizes the total weighted overlap is called *optimal*.

The formulation our *Weighted Single-Processor Scheduling* problem (`WSPS`), is as follows.

Instance: A set of weighted jobs $\mathcal{J}$ with arbitrary given processing times.

Goal: Find an optimal schedule for $\mathcal{J}$.

4 Preliminaries

Let $\mathcal{S}$ be a feasible schedule. We denote by $s_{\mathcal{S}}^{\mathcal{M}}(j)$ and $C_{\mathcal{S}}^{\mathcal{M}}(j)$ the start time and the completion times of a job j on the shared processor, respectively. For brevity we take $s_{\mathcal{S}}^{\mathcal{M}}(j) = C_{\mathcal{S}}^{\mathcal{M}}(j) = 0$ if a job j executes on its private processor only. Whenever $s_{\mathcal{S}}^{\mathcal{M}}(j) < C_{\mathcal{S}}^{\mathcal{M}}(j)$, i.e., some non-empty part of a job j executes on $\mathcal{M}$, then we say that the job j *appears on* $\mathcal{M}$ in schedule $\mathcal{S}$. If, in a schedule $\mathcal{S}$, there is no idle time on the shared processor in time interval

$$\left[0, \max\{C_{\mathcal{S}}^{\mathcal{M}}(j) \;\middle|\; j \in \mathcal{J}\}\right],$$

then we say that $\mathcal{S}$ *has no gaps*. We have the following results form the literature.

Observation 4.1 ([6]) *There exists an optimal schedule that has no gaps.*

A schedule $\mathcal{S}$ is called *non-preemptive* if each job j executes in $\mathcal{S}$ in time interval $[s^{\mathcal{M}}_{\mathcal{S}}(j), C^{\mathcal{M}}_{\mathcal{S}}(j)]$ on the shared processor. We say that a schedule is *synchronized* if it satisfies the following conditions:

(i) it is non-preemptive and has no gaps,

(ii) for each job j that appears on the shared processor it holds $C^{\mathcal{M}}_{\mathcal{S}}(j) = C^{\mathcal{P}}_{\mathcal{S}}(j)$.

Theorem 1 ([6]) *There exists an optimal synchronized schedule.*

Consider a synchronized schedule $\mathcal{S}$. Let $A = \{j_1, \ldots, j_k\} \subseteq \mathcal{J}$ be the set of all jobs that appear on $\mathcal{M}$ in $\mathcal{S}$, where the jobs are ordered according to their completion times in $\mathcal{S}$, i.e. $C^{\mathcal{M}}_{\mathcal{S}}(j_1) < \cdots < C^{\mathcal{M}}_{\mathcal{S}}(j_k)$. Note that the set A and the order are enough to determine the schedule $\mathcal{S}$. Indeed, given the order $(j_1, \ldots, j_k)$ we obtain that for each $i \in \{1, \ldots, k\}$ (by proceeding with increasing values of i),

$$s^{\mathcal{M}}_{\mathcal{S}}(j_i) = C^{\mathcal{M}}_{\mathcal{S}}(j_{i-1}) \quad \text{and} \quad C^{\mathcal{M}}_{\mathcal{S}}(j_i) = C^{\mathcal{P}}_{\mathcal{S}}(j_i) = \left(p_i + s^{\mathcal{M}}_{\mathcal{S}}(j_i)\right)/2,$$

where $C^{\mathcal{M}}_{\mathcal{S}}(j_0) = 0$. This formula implies that the start times and completion times can be iteratively computed for all jobs. Thus for synchronized schedules we write for brevity $\mathcal{S} = (j_1, \ldots, j_k)$ to refer to the schedule computed above.

5 A $O(nlogn)$-time optimal algorithm for antithetical instances

We call an instance $\mathcal{J}$ of the problem *antithetical* if for any two jobs i and j it holds: $p_i \leq p_j$ implies $w_i \geq w_j$. We call a schedule $\mathcal{S}$ *processing time ordered* if $\mathcal{S} = (j_1, \ldots, j_n)$, where $p_{j_i} \leq p_{j_{i+1}}$ for each $i \in \{1, \ldots, n-1\}$. Observe that by construction, $\mathcal{S}$ is synchronized and all jobs from $\mathcal{J}$ appear on the shared processor, see [10] and [7]. We now prove that any processing time ordered schedule is an optimal solution for an antithetical instance. This gives an $O(n \log n)$-time optimization algorithm for antithetical instances. We remark that this algorithm generalizes the previously known solutions for the unweighted case ($w_1 = \cdots = w_n$) from [7, 10].

Lemma 5.1 *A processing time ordered schedule is optimal for any antithetical instance of the problem* `WSPS`.

Proof: Let $\mathcal{S}$ be an optimal schedule for an antithetical instance $\mathcal{J}$. By Theorem 1 we can assume that $\mathcal{S}$ is synchronized. We assume without loss of generality that the jobs in $\mathcal{J} = \{1, \ldots, n\}$ are ordered in non-decreasing order of their processing times, i.e. $p_1 \leq p_2 \leq \cdots \leq p_n$. Let $A \subseteq \mathcal{J}$ be the set of jobs that appear on the shared processor in $\mathcal{S}$. Let $\pi(1), \pi(2), \ldots, \pi(k)$, $k = |A|$, be the order of jobs on the shared processor in $\mathcal{S}$, i.e., $\mathcal{S} = (\pi(1), \pi(2), \ldots, \pi(k))$.

The largest index $i \in \{1, \ldots, k-1\}$ such that either

$$p_{\pi(i)} > p_{\pi(i+1)} \tag{1}$$

or

$$\text{there exists } \bar{j} \in \mathcal{J} \setminus A \text{ such that } \bar{j} > \pi(i+1) \tag{2}$$

is called the *violation point of* $\mathcal{S}$; set the violation point to be zero if no such index i exists. Note that violation point equals zero only for processing time ordered schedules.

Among all optimal and synchronized schedules we take $\mathcal{S}$ to satisfy the following:

(a) the number of jobs that appear on $\mathcal{M}$ in $\mathcal{S}$ is maximum, and

(b) with respect to (a): the violation point of $\mathcal{S}$ is minimum.

We aim at proving that the violation point of $\mathcal{S}$ is zero, that is:

$$A = \mathcal{J} \tag{3}$$

and

$$p_{\pi(i)} \leq p_{\pi(i+1)} \text{ for each } i \in \{1, \ldots, k-1\}, \tag{4}$$

which immediately implies the lemma.

We prove these claims by contradiction. Let $i > 0$ be the violation point of $\mathcal{S}$. By definition, we have that one of the cases (1) or (2) holds. We should arrive at a contradiction in both cases and we start by analyzing the case of (1), that is, we assume that (1) holds for the violation point i. For antithetical instances we have $w_{\pi(i)} \leq w_{\pi(i+1)}$. Also, for convenience let without loss of generality we denote $j = \pi(i+1)$ and $j+1 = \pi(i)$. Thus $p_{j+1} > p_j$ by (1). In the next two paragraphs we describe a transition from $\mathcal{S}$ to a new schedule $\mathcal{S}'$. This transition is depicted in Figure 1.

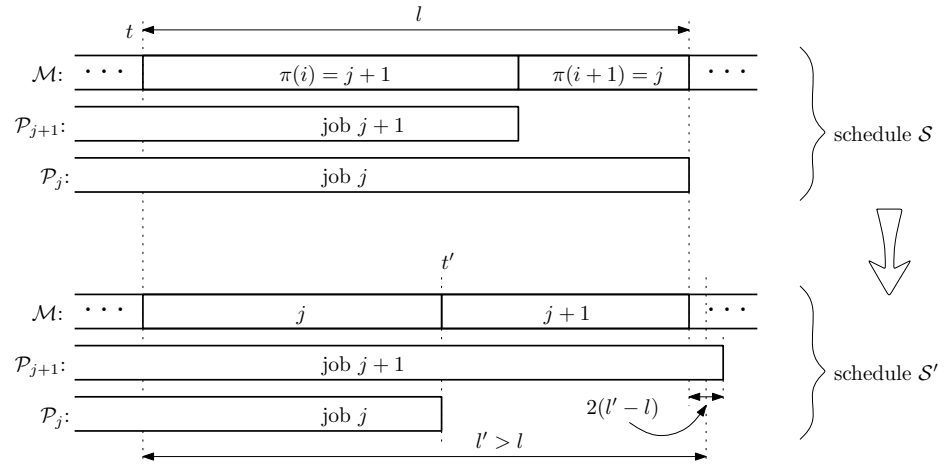

Figure 1: Transition from $\mathcal{S}$ to $\mathcal{S}'$ in proof of Lemma 5.1

Consider the intervals in which the two jobs j and $j+1$ execute on $\mathcal{M}$ in $\mathcal{S}$ and suppose that $j+1$ starts at t on $\mathcal{M}$, $s_{\mathcal{S}}^{\mathcal{M}}(j+1) = t$. Since $\mathcal{S}$ is synchronized, the length $l = C_{\mathcal{S}}^{\mathcal{M}}(j) - s_{\mathcal{S}}^{\mathcal{M}}(j+1)$ of this sequence on the shared processor equals

$$l = \left(C_{\mathcal{S}}^{\mathcal{M}}(j) - s_{\mathcal{S}}^{\mathcal{M}}(j)\right) + \left(C_{\mathcal{S}}^{\mathcal{M}}(j+1) - s_{\mathcal{S}}^{\mathcal{M}}(j+1)\right) = \frac{p_j - t}{2} + \frac{p_{j+1} - t}{4}, \tag{5}$$

and its contribution x to the value of objective function (total weighted overlap) equals

$$x = \left(C_{\mathcal{S}}^{\mathcal{M}}(j+1) - s_{\mathcal{S}}^{\mathcal{M}}(j+1)\right) w_{j+1} + \left(C_{\mathcal{S}}^{\mathcal{M}}(j) - s_{\mathcal{S}}^{\mathcal{M}}(j)\right) w_j = \frac{p_{j+1} - t}{2}\left(w_{j+1} - \frac{w_j}{2}\right) + \frac{p_j - t}{2} w_j. \tag{6}$$

Thus, we can express the total weighted overlap of $\mathcal{S}$ as follows:

$$\Sigma(\mathcal{S}) = c + x \quad \text{for some } c \in \mathbb{R}. \tag{7}$$

Before we formally define $\mathcal{S}'$, we analyze the impact the reversed order of the two jobs j and $j+1$ in the interval $(t, t+l)$ has on $\mathcal{S}$ and its objective function. Suppose for the time being that j starts at t on $\mathcal{M}$ and is followed by the job $j+1$ and that either job is executed in such a way that it completes on both $\mathcal{M}$ and its private processors at the same time. Then the length l' of the interval $(t, t+l')$ occupied by these two jobs on the shared processor equals

$$l' = \frac{p_{j+1} - t}{2} + \frac{p_j - t}{4}$$

and its contribution x' to the value of objective function equals

$$x' = \frac{p_j - t}{2}\left(w_j - \frac{w_{j+1}}{2}\right) + \frac{p_{j+1} - t}{2} w_{j+1}. \tag{8}$$

Clearly, $l' > l$ for $p_{j+1} > p_j$ by

$$l' - l = \frac{p_{j+1} - p_j}{4}. \tag{9}$$

The job j completes at

$$t' = t + \frac{p_j - t}{2} < t + \frac{p_j - t}{2} + \frac{p_{j+1} - t}{4} = t + l \tag{10}$$

after the exchange. We construct the schedule $\mathcal{S}'$ as follows: $\mathcal{S}$ and $\mathcal{S}'$ are identical in time intervals $[0, t)$ and $(t+l, +\infty)$, the job j executes in time interval

$$(t, t') = (t, t + (p_j - t)/2) \tag{11}$$

in $\mathcal{S}'$ and the job $j+1$ executes in time interval

$$(t', t+l) = (t + (p_j - t)/2, t + l)$$

on processor $\mathcal{M}$ in $\mathcal{S}'$. Note that j finishes at the same time on its private and shared processor in $\mathcal{S}'$ while the job $j+1$ does not have this property. Since $p_{j+1} > p_j$, $j+1$ completes $(p_{j+1} - p_j)/4$ units later on its private processor. Thus $\mathcal{S}'$ is not synchronized; see also Figure 1. The total weighted overlap of $\mathcal{S}'$ is then by (5), (10) and (11):

$$\begin{aligned} \Sigma(\mathcal{S}') = &\quad c + (t' - t)w_j + (t + l - t')w_{j+1}, \\ = &\quad c + \frac{p_j - t}{2} w_j + \frac{p_{j+1} - t}{4} w_{j+1} \end{aligned} \tag{12}$$

where c is defined in (7). By (7) and (12) we obtain that the difference (in total weighted overlaps) between $\mathcal{S}'$ and $\mathcal{S}$ is

$$\Sigma(\mathcal{S}') - \Sigma(\mathcal{S}) = \frac{p_{j+1} - t}{4}(w_j - w_{j+1}). \tag{13}$$

Since $p_{j+1} - t > 0$ (this holds since the job $j+1$ appears on the shared processor in $\mathcal{S}$) and $w_j \geq w_{j+1}$, we have that $\Sigma(\mathcal{S}') - \Sigma(\mathcal{S}) \geq 0$. Note that if w_j is strictly greater than w_{j+1}, then we obtain the desired contradiction with the optimality of $\mathcal{S}$. However, if $w_j = w_{j+1}$, or in other words $\Sigma(\mathcal{S}) = \Sigma(\mathcal{S}')$, then we need to use different arguments to arrive at a contradiction.

To that end denote

$$q = \sum_{i'=i+2}^{k} \frac{w_{\pi(i')}}{2^{i'-i-1}}.$$

We show that

$$w_{j+1} \geq q. \tag{14}$$

Suppose for a contradiction that $w_{j+1} < q$. To obtain a contradiction with this assumption we will convert $\mathcal{S}'$ into a schedule $\mathcal{S}''$ with strictly greater total weighted overlap, which will contradict the optimality of the original schedule $\mathcal{S}$. This conversion is described in the next two paragraphs and depicted in Figure 2. Let

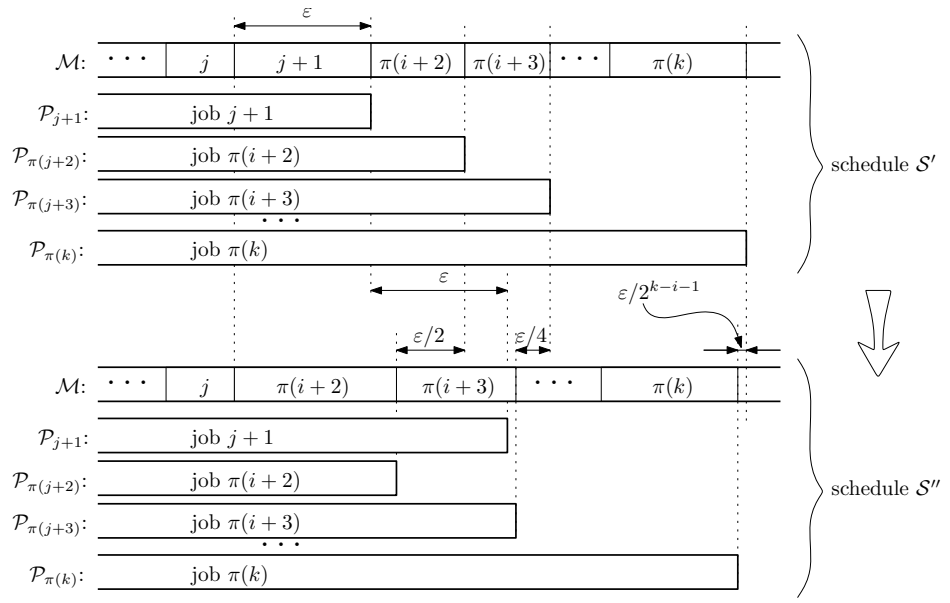

Figure 2: Transition from $\mathcal{S}'$ to $\mathcal{S}''$ in proof of Lemma 5.1 when proving (14)

$$\varepsilon = C^{\mathcal{M}}_{\mathcal{S}'}(j+1) - s^{\mathcal{M}}_{\mathcal{S}'}(j+1).$$

Observe that $\varepsilon > 0$ by (9). By assumption $\mathcal{S}$ is synchronized and $\mathcal{S}$ and $\mathcal{S}'$ are identical on $\mathcal{M}$ in time intervals $[C^{\mathcal{M}}_{\mathcal{S}'}(j+1), +\infty)$. Thus

$$C^{\mathcal{M}}_{\mathcal{S}'}(i') = C^{\mathcal{M}}_{\mathcal{S}}(i') = C^{\mathcal{P}}_{\mathcal{S}}(i') = C^{\mathcal{P}}_{\mathcal{S}'}(i')$$

for each job i' that appears on $\mathcal{M}$ and completes in $\mathcal{S}'$ later than the job $j+1$. The schedule $\mathcal{S}''$ is defined as follows. Let, $\mathcal{S}''$ and $\mathcal{S}'$ be identical on $\mathcal{M}$ in time interval

$$\left[0, C^{\mathcal{M}}_{\mathcal{S}'}(j+1) - \varepsilon\right) = \left[0, s^{\mathcal{M}}_{\mathcal{S}'}(j+1)\right).$$

Then, the job $j+1$ is not present on $\mathcal{M}$ in $\mathcal{S}'$. It executes only on $\mathcal{P}_{j+1}$ in $\mathcal{S}''$. Finally, for each job $\pi(i')$, $i' \in \{i+2, \ldots, k\}$, we set:

$$s^{\mathcal{M}}_{\mathcal{S}''}(\pi(i')) = s^{\mathcal{M}}_{\mathcal{S}'}(\pi(i')) - \frac{\varepsilon}{2^{i'-i-2}},$$

$$C^{\mathcal{M}}_{\mathcal{S}''}(\pi(i')) = C^{\mathcal{M}}_{\mathcal{S}'}(\pi(i')) - \frac{\varepsilon}{2^{i'-i-1}},$$

$$C^{\mathcal{P}}_{\mathcal{S}''}(\pi(i')) = C^{\mathcal{P}}_{\mathcal{S}'}(\pi(i')) - \frac{\varepsilon}{2^{i'-i-1}}.$$

Both $\mathcal{S}''$ and $\mathcal{S}'$ are the same on $\mathcal{P}_{i'}$ for each $i' \in \mathcal{J} \setminus \{j+1, \pi(i+2), \ldots, \pi(k)\}$, i.e., on each processor not specified by the formulas above.

Clearly, $\mathcal{S}''$ is feasible and synchronized. To compare its total weighted overlap to that of $\mathcal{S}'$, note that on the one hand the value of $\mathcal{S}''$ decreases by εw_{j+1} in comparison to $\mathcal{S}'$ since the job $j+1$ does not appear on $\mathcal{M}$ in $\mathcal{S}''$, on the other hand it increases since the subintervals with the jobs that follow $j+1$ on $\mathcal{M}$ get longer due to $j+1$ disappearance from $\mathcal{M}$. Hence

$$\Sigma(\mathcal{S}'') = \Sigma(\mathcal{S}') - \varepsilon w_j + \varepsilon \sum_{i'=i+2}^{k} \frac{w_{i'}}{2^{i'-i-1}} = \Sigma(\mathcal{S}') + \varepsilon(q - w_j) > 0$$

because $\varepsilon > 0$ and $w_j < q$ by assumption. Thus, we obtain a contradiction with the optimality of $\mathcal{S}$ (recall that $\Sigma(\mathcal{S}) \leq \Sigma(\mathcal{S}')$).

Having proved (14), we complete the proof of case (1) by performing another transformation of the schedule $\mathcal{S}'$ to a new schedule $\mathcal{S}''$. We intend this transformation to also apply to the case of (2). Hence we will conduct the remaining part of the proof in such a way that we complete the proof of case (1) and carry out a complete proof of case (2) all at one time. Since, we will obtain a contradiction in both cases, the proof of both (3) and (4) will be completed. Thus, what we need is to define the $\mathcal{S}'$ and the job $j+1$ in order to include the case (2) into our proof: for (2) we take $\mathcal{S}' := \mathcal{S}$, $j = \pi(i+1)$ and $j+1 = \bar{j}$. (Recall that $\bar{j}$ is defined in (2) to be a job that does not appear on $\mathcal{M}$ and has processing time that is greater or equal to that of $\pi(i+1)$.) It will follow from the construction below that the transformation works for both $\mathcal{S}'$, j and $j+1$ used in the earlier part of the proof of case (1) and for the new $\mathcal{S}$, $\pi(i)$ and $\bar{j}$ in (2). Observe that, informally speaking, the only difference in both cases is that in case (1) the job $j+1$ is present on $\mathcal{M}$ and completes on $\mathcal{P}$ later than on $\mathcal{M}$ while in case (2) the job $j+1$ is not present on $\mathcal{M}$. We have in particular

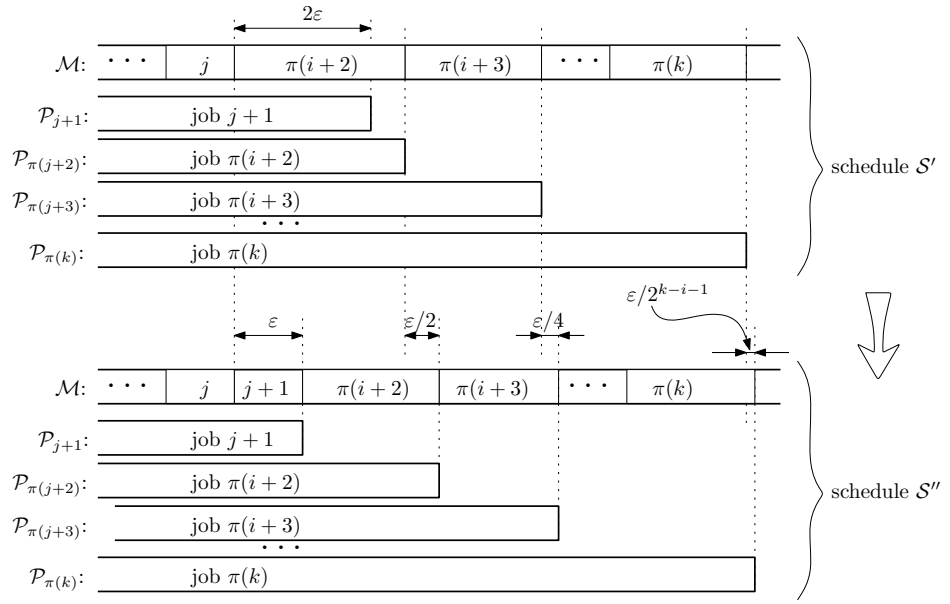

Figure 3: Transition from $\mathcal{S}'$ to $\mathcal{S}''$ in proof of Lemma 5.1 for the completion of the proof of (3) and (4); in this particular example we have that the job $j+1$ does not appear on $\mathcal{M}$, thus presenting the case when we take $j+1 = \bar{j}$

$$C^{\mathcal{P}}_{\mathcal{S}'}(j+1) > C^{\mathcal{M}}_{\mathcal{S}'}(j+1). \tag{15}$$

Let

$$\varepsilon = \frac{1}{2}\left(C^{\mathcal{P}}_{\mathcal{S}'}(j+1) - s^{\mathcal{M}}_{\mathcal{S}'}(\pi(i+2))\right).$$

We now obtain a schedule $\mathcal{S}''$ as follows. First, $\mathcal{S}''$ and $\mathcal{S}'$ are identical on $\mathcal{M}$ in time interval $[0, C^{\mathcal{M}}_{\mathcal{S}'}(j))$,

$$s^{\mathcal{M}}_{\mathcal{S}''}(\pi(i')) = s^{\mathcal{M}}_{\mathcal{S}'}(\pi(i')) + \frac{\varepsilon}{2^{i'-i-2}}$$

and

$$C^{\mathcal{M}}_{\mathcal{S}''}(\pi(i')) = C^{\mathcal{P}}_{\mathcal{S}''}(\pi(i')) = C^{\mathcal{M}}_{\mathcal{S}'}(\pi(i')) + \frac{\varepsilon}{2^{i'-i-1}}$$

for each $i' \in \{i+2, \ldots, k\}$. Informally, the start of the $\pi(i')$-th job is postponed on $\mathcal{M}$ by $\varepsilon/2^{i'-i-2}$ and its completion time is set up in such a way that this job completes on both $\mathcal{M}$ and $\mathcal{P}$ at the same time. Then, $s^{\mathcal{M}}_{\mathcal{S}''}(j+1) = C^{\mathcal{M}}_{\mathcal{S}''}(j)$, $C^{\mathcal{M}}_{\mathcal{S}''}(j+1) = s^{\mathcal{M}}_{\mathcal{S}''}(\pi(i+2))$ and the remaining part of $j+1$ executes on $\mathcal{P}_{j+1}$. Finally, the execution of the remaining jobs (those for which completion times on their private processors have not been defined above) on their private processors remains the same in $\mathcal{S}''$ as in $\mathcal{S}'$.

We argue that the schedule $\mathcal{S}''$ is feasible. Clearly, the new time intervals in which job $j+1$ executes on $\mathcal{M}$ and $\mathcal{P}_{j+1}$ are correct. Moreover, $j+1$ completes on both processors at the same time. For the jobs that follow $j+1$ on $\mathcal{M}$, each of them also completes at the same time on $\mathcal{M}$ and its private processor. Therefore these jobs are executed correctly because by the choice of i, they are ordered on $\mathcal{M}$ (in $\mathcal{S}'$ and also on $\mathcal{S}''$) with non-decreasing values of their processing times:

$$p_{j+1} \leq p_{\pi(i+2)} \leq \cdots \leq p_{\pi(k)}.$$

This implies that $\mathcal{S}''$ is feasible and synchronized.

From the construction we obtain that

$$\Sigma(\mathcal{S}'') = \Sigma(\mathcal{S}') + \varepsilon w_{j+1} - \varepsilon \sum_{i'=i+2}^{k} \frac{w_{\pi(i')}}{2^{i'-i-1}}. \tag{16}$$

Therefore, in case (1), we obtain by (14) that $\Sigma(\mathcal{S}'') \geq \Sigma(\mathcal{S}')$. In particular, if $\Sigma(\mathcal{S}'') > \Sigma(\mathcal{S}')$, then we have immediately a contradiction with the optimality of $\mathcal{S}$ since $\Sigma(\mathcal{S}') \geq \Sigma(\mathcal{S})$. On the other hand, if $\Sigma(\mathcal{S}'') = \Sigma(\mathcal{S}')$, then the contradiction comes from the selection of $\mathcal{S}$ to be a schedule that minimizes the violation point i (both $\mathcal{S}''$ and $\mathcal{S}$ have the same number of jobs present on the shared processor).

To consider case (2), observe that since i is the violation point of $\mathcal{S}$, $w_{\pi(i+2)} \geq w_{\pi(i')}$ for each $i' \in \{i+2, \ldots, k\}$. Therefore,

$$\sum_{i'=i+2}^{k} \frac{w_{\pi(i')}}{2^{i'-i-1}} \leq w_{\pi(i+2)} \sum_{i'=i+2}^{k} \frac{1}{2^{i'-i-1}} < w_{\pi(i+2)}.$$

Hence by (16), $\Sigma(\mathcal{S}'') > \Sigma(\mathcal{S}')$ because $w_{j+1} \geq w_{\pi(i+2)}$. The latter inequality comes from the fact that the instance is antithetical and $p_{j+1} \leq p_{\pi(i+2)}$, which comes from the maximality of i with respect to the conditions (1) and (2). Thus, we again have a contradiction with the choice of $\mathcal{S}$, which completes the proof of the lemma. □

6 A 1/2-approximation Algorithm

Let $0 = q_0 < q_1 < \cdots < q_\ell$ and $u_1, \ldots, u_\ell \geq 0$ for some $\ell \geq 1$. An *envelope* for $q_1, \ldots, q_\ell$ and $u_1, \ldots, u_\ell$ is a step-function of non-negative x defined as follows

$$e(q_1, \ldots, q_\ell, u_1, \ldots, u_\ell, x) = \begin{cases} u_1 & \text{if } q_0 \leq x \leq q_1 \\ u_2 & \text{if } q_1 < x \leq q_2 \\ & \cdots \\ u_\ell & \text{if } q_{k-1} < x \leq q_\ell \\ 0 & \text{if } q_\ell < x. \end{cases}$$

The *area* of the envelope e is

$$\sum_{i=1}^{\ell} u_i(q_i - q_{i-1}).$$

Let $\mathcal{J}$ be a set of jobs. Without loss of generality we assume $p_1 \leq \cdots \leq p_n$, for any tie we assume the jobs that are tied are ordered in ascending order of their weights, i.e. the heaviest tied job comes last in the tie. A sequence of jobs $i_1, \ldots, i_\ell$ for some $\ell \geq 1$, where $1 \leq i_1 < \cdots < i_\ell \leq n$, is called a *key sequence* for $\mathcal{J}$ if it satisfies the following conditions:

(i) $i_\ell = n$,

(ii) $w_{i_1} > \cdots > w_{i_\ell}$,

(iii) $w_k \leq w_{i_j}$ for each $k \in I_{i_j} = \{i_{j-1} + 1, \ldots, i_j\}$ and $j \in \{1, \ldots, \ell\}$, where $i_0 = 0$.

Clearly $p_{i_1} < \cdots < p_{i_\ell}$, thus $u(p_{i_1}, \ldots, p_{i_\ell}, w_{i_1}, \ldots, w_{i_\ell}, x)$ is an envelope; we refer to it as the *upper envelope* for $\mathcal{J}$. Let u^* be the area of the upper envelope for $\mathcal{J}$.

Note that the key sequence always exists. This follows from the fact that it can be constructed 'greedily' by starting with picking the last job of the sequence (see (i)) and then iteratively selecting the predecessor of the previously selected job so that the predecessor has strictly bigger weight (see (ii)) and satisfies the condition (iii). Also, the key sequence is unique by the same argument.

We have the following simple observation.

Claim 6.1 *For each $k \in \{1, \ldots, \ell\}$, $w_{i_k} = \max\{w_j \mid i_k \leq j \leq n\}$.* □

The key sequence for $\mathcal{J}$ defines a synchronized schedule $\mathcal{S}_{\text{key}}$ for $\mathcal{J}$ with the set of jobs executed on the shared processor being $\mathcal{J}_{\text{key}} = \{i_1, \ldots, i_\ell\}$ and the permutation of the jobs on the processor being $\pi(j) = i_j$ for $j \in \{1, \ldots, \ell\}$. The jobs in $\mathcal{J} \setminus \mathcal{J}_{\text{key}}$ are executed on their private processors only. Following our notation introduced in Section 4, we get $\mathcal{S}_{\text{key}} = (i_1, \ldots, i_\ell)$. We have the following lemma.

Lemma 6.2 *For the schedule $\mathcal{S}_{\text{key}}$ it holds $2\Sigma(\mathcal{S}_{\text{key}}) \geq u^*$.*

Proof: We argue that for each $k \in \{1, \ldots, \ell\}$,

$$w_{i_k}\left(C^{\mathcal{M}}_{\mathcal{S}_{\text{key}}}(i_k) - s^{\mathcal{M}}_{\mathcal{S}_{\text{key}}}(i_k)\right) \geq w_{i_k}\frac{p_{i_k} - p_{i_{k-1}}}{2}, \tag{17}$$

where $p_{i_0} = 0$. Note that $s^{\mathcal{M}}_{\mathcal{S}_{\rm key}}(i_k) \le p_{i_{k-1}}$ for each $k \in \{1, \ldots, \ell\}$. Thus,

$$C^{\mathcal{M}}_{\mathcal{S}_{\rm key}}(i_k) - s^{\mathcal{M}}_{\mathcal{S}_{\rm key}}(i_k) = \frac{s^{\mathcal{M}}_{\mathcal{S}_{\rm key}}(i_k) + p_{i_k}}{2} - s^{\mathcal{M}}_{\mathcal{S}_{\rm key}}(i_k) \ge \frac{p_{i_k} - p_{i_{k-1}}}{2}$$

for each $k \in \{1, \ldots, \ell\}$, which proves (17).

By (17),

$$\Sigma(\mathcal{S}_{\rm key}) = \sum_{k=1}^{\ell} w_{i_k} \left(C^{\mathcal{M}}_{\mathcal{S}_{\rm key}}(i_k) - s^{\mathcal{M}}_{\mathcal{S}_{\rm key}}(i_k) \right) \ge \sum_{k=1}^{\ell} w_{i_k} \frac{p_{i_k} - p_{i_{k-1}}}{2} = \frac{u^*}{2}.$$

□

We now prove that the area u^* of the upper envelope for $\mathcal{J}$ is an upper bound on the value of optimal solution for $\mathcal{J}$. By Theorem 1, there exists an optimal synchronized schedule $\mathcal{S}_{\rm opt}$ for $\mathcal{J}$. Let $\mathcal{J}_{\rm opt} \subseteq \mathcal{J}$ be the set of jobs that appear on the shared processor in $\mathcal{S}_{\rm opt}$ and let π be a permutation of jobs in $\mathcal{J}_{\rm opt}$ in $\mathcal{S}_{\rm opt}$. Thus, we have $\mathcal{S}_{\rm opt} = (\pi(1), \ldots, \pi(|\mathcal{J}_{\rm opt}|))$. It holds $0 < C^{\mathcal{M}}_{\mathcal{S}_{\rm opt}}(\pi(1)) < \cdots < C^{\mathcal{M}}_{\mathcal{S}_{\rm opt}}(\pi(|\mathcal{J}_{\rm opt}|)) < p_n$ and therefore

$$e\left(C^{\mathcal{M}}_{\mathcal{S}_{\rm opt}}(\pi(1)), \ldots, C^{\mathcal{M}}_{\mathcal{S}_{\rm opt}}(\pi(|\mathcal{J}_{\rm opt}|)), w_{\pi(1)}, \ldots, w_{\pi(|\mathcal{J}_{\rm opt}|)}, x \right).$$

is an envelope for $C^{\mathcal{M}}_{\mathcal{S}_{\rm opt}}(\pi(1)), \ldots, C^{\mathcal{M}}_{\mathcal{S}_{\rm opt}}(\pi(|\mathcal{J}_{\rm opt}|))$ and $w_{\pi(1)}, \ldots, w_{\pi(|\mathcal{J}_{\rm opt}|)}$. Let the area of this envelope be e^*. We have the following key result.

Lemma 6.3 *It holds $e^* \le u^*$.*

Proof: Observe that for each index $j \in \{1, \ldots, |\mathcal{J}_{\rm opt}|\}$ there exists $\tau(j) \in \{1, \ldots, n\}$ such that

$$C^{\mathcal{M}}_{\mathcal{S}_{\rm opt}}(\pi(j)) \le p_{i_{\tau(j)}}. \tag{18}$$

This follows from condition (i) in definition of key sequence. If, for a given j, there are several jobs $\tau(j)$ that satisfy the above, then take $\tau(j)$ to be the minimum one.

We argue that

$$w_{\pi(j)} \le w_{i_{\tau(j)}} \tag{19}$$

for each $j \in \{1, \ldots, |\mathcal{J}_{\rm opt}|\}$. Suppose for a contradiction that (19) does not hold. We consider two cases. In the first case let

$$p_{\pi(j)} \le p_{i_{\tau(j)}}.$$

By condition (iii) in definition of the key sequence and the minimality of $\tau(j)$, $\pi(j)$ does not belong to the key sequence. But then, $w_{\pi(j)} > w_{i_{\tau(j)}}$ implies that there is t such that $p_{\pi(j)} < p_{i_t} < p_{i_{\tau(j)}}$, which contradicts the choice of $\tau(j)$. In the second case let

$$p_{\pi(j)} > p_{i_{\tau(j)}}.$$

Take the minimum index t such that $p_{i_t} \ge p_{\pi(j)}$. By condition (iii) in definition of the key sequence, $w_{i_t} \ge w_{\pi(j)}$. Since $w_{\pi(j)} > w_{i_{\tau(j)}}$, condition (ii) in definition of the key sequence implies that $i_{\tau(j)}$ does not belong to the key sequence — a contradiction. This completes the proof of (19).

Since the upper envelope for $\mathcal{J}$ is non-increasing function in x, we obtain by (18) and (19) that

$$e\left(C^{\mathcal{M}}_{\mathcal{S}_{\rm opt}}(\pi(1)), \ldots, C^{\mathcal{M}}_{\mathcal{S}_{\rm opt}}(\pi(|\mathcal{J}_{\rm opt}|)), w_{\pi(1)}, \ldots, w_{\pi(|\mathcal{J}_{\rm opt}|)}, x \right) \le u\left(p_{i_1}, \ldots, p_{i_\tau}, w_{i_1}, \ldots, w_{i_\tau}, x\right)$$

for each $x \ge 0$, which completes the proof. □

Since $\Sigma(\mathcal{S}_{\rm opt}) = e^*$ and $\Sigma(\mathcal{S}_{\rm key}) \le \Sigma(\mathcal{S}_{\rm opt})$, Lemmas 6.2 and 6.3 give the following.

Corollary 2 *It holds* $e^*/2 \leq \Sigma(\mathcal{S}_{\text{key}}) \leq e^*$. □

Theorem 3 *The key sequence for* $\mathcal{J}$ *provides a* $1/2$*-approximation solution to the problem* WSPS*. This sequence can be found in time* $O(n \log n)$ *for any set of jobs* $\mathcal{J}$*, where* $n = |\mathcal{J}|$*. Moreover, the bound of* $1/2$ *is tight, i.e., for each* $\varepsilon > 0$ *there exists a problem instance such that* $\Sigma(\mathcal{S}_{\text{key}}) < \left(\frac{1}{2} + \varepsilon\right)\Sigma(\mathcal{S}_{\text{opt}})$.

Proof: The fact that the key sequence is a $1/2$-approximation of the optimal solution follows from Corollary 2. The key sequence can be constructed directly from the definition and sorting the jobs in $\mathcal{J}$ according to their processing times determines the $O(n \log n)$ running time.

To close we show that the $1/2$ bound for the key sequences is tight. Take $\mathcal{J}$ to contain n jobs, each of the same weight $w > 0$ and the same length $p > 0$. The key sequence consists of one job and therefore the value of the corresponding schedule $\mathcal{S}_{\text{key}}$ is $\Sigma(\mathcal{S}_{\text{key}}) = wp/2$. Take a schedule $\mathcal{S}$ that places all jobs in $\mathcal{J}$ on the shared processor. We have $\Sigma(\mathcal{S}) = wp(1 - 1/2^n)$. For $\varepsilon > 0$ take $n = \lceil \log(1/\varepsilon) \rceil$. If $\mathcal{S}_{\text{opt}}$ is an optimal schedule, then

$$\frac{\Sigma(\mathcal{S}_{\text{key}})}{\Sigma(\mathcal{S}_{\text{opt}})} \leq \frac{\Sigma(\mathcal{S}_{\text{key}})}{\Sigma(\mathcal{S})} = \frac{1}{2} - \frac{1}{2^{n+1}} \leq \frac{1}{2} - \varepsilon.$$

□

7 Open problems and further research

The complexity status of WSPS remains open. The generalized problem with multiple shared processors is strongly NP-hard [6] when the number of shared processors is a part of the input. However, it remains open whether the generalized problem with fixed number of processors is NP-hard (i.e., whether the multiprocessor problem is FPT with respect to the number of processors). This complexity result and the open complexity questions clearly underline the difficulty in finding efficient optimization algorithms for the shared processor scheduling problem. The development of an efficient branch-and-bound algorithm for the problem remains unexplored so far. The 1/2-approximation algorithm along with the structural properties of optimal schedules presented in this paper and in [6] may prove useful building blocks of such an algorithm.

Acknowledgements

This research has been supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Grant OPG0105675 and by Polish National Science Center under contract DEC-2011/02/A/ST6/00201.

References

[1] A. Agnetis, J.-C. Billaut, S. Gawiejnowicz, D. Pacciarelli, and A. Soukhal. *Multiagent Scheduling. Models and Algorithms*. Springer, 2014.

[2] E. J. Anderson. A new continuous model for job–shop scheduling. *International Journal of System Science*, 12:1469–1475, 1981.

[3] V. Bharadwaj, D. Ghose, and T.G. Robertazzi. Divisible load theory: A new paradigm for load scheduling in distributed systems. *Cluster Computing*, 6:7–17, 2003.

[4] Y. Bukchin and E. Hanany. Decentralization cost in scheduling: A game-theoretic approach. *Manufacturing & Service Operations Management*, 9(3):263–275, 2007.

[5] T. C. E. Cheng, C. T. Ng, and J. J. Yuan. Multi-agent scheduling on a single machine to minimize total weighted number of tardy jobs. *Theor. Comput. Sci.*, 362(1-3):273–281, 2006.

[6] D. Dereniowski and W. Kubiak. Shared multi-processor scheduling. *CoRR*, abs/1607.06021, 2016.

[7] B. Hezarkhani and W. Kubiak. Decentralized subcontractor scheduling with divisible jobs. *J. Scheduling*, 18(5):497–511, 2015.

[8] K. Lee, B.-C. Choi, J. Y.-T. Leung, and M. L. Pinedo. Approximation algorithms for multi-agent scheduling to minimize total weighted completion time. *Inf. Process. Lett.*, 109(16):913–917, 2009.

[9] G. L. Vairaktarakis. Noncooperative games for subcontracting operations. *Manufacturing and Service Operations Management*, 15:148–158, 2013.

[10] G.L. Vairaktarakis and T. Aydinliyim. Centralization versus competition in subcontracting operations. Technical Memorandum Number 819, Case Western Reserve University, 2007.